

INVERTING THE FURSTENBERG CORRESPONDENCE

JEREMY AVIGAD

ABSTRACT. Given a sequence of sets $A_n \subseteq \{0, \dots, n-1\}$, the Furstenberg correspondence principle provides a shift-invariant measure on $2^{\mathbb{N}}$ that encodes combinatorial information about infinitely many of the A_n 's. Here it is shown that this process can be inverted, so that for any such measure, ergodic or not, there are finite sets whose combinatorial properties approximate it arbitrarily well. The finite approximations are obtained from the measure by an explicit construction, with an explicit upper bound on how large n has to be to yield a sufficiently good approximation.

We draw conclusions for computable measure theory, and show, in particular, that given any computable shift-invariant measure on $2^{\mathbb{N}}$, there is a computable element of $2^{\mathbb{N}}$ that is generic for the measure. We also consider a generalization of the correspondence principle to countable discrete amenable groups, and once again provide an effective inverse.

1. INTRODUCTION

For each n , let A_n be a subset of $\{0, \dots, n-1\}$. The Furstenberg correspondence principle, described more precisely in Section 2, allows one to assign a shift-invariant measure on Cantor space, $2^{\mathbb{N}}$, which encodes combinatorial information about infinitely many of the A_n 's. This correspondence lies at the heart of Furstenberg's remarkable ergodic-theoretic proof [10, 11] of Szemerédi's theorem, and allows one to use facts about shift-invariant measures on Cantor space to draw conclusions about subsets of $\{0, \dots, n-1\}$, for sufficiently large n .

It is natural to ask exactly which shift-invariant measures on $2^{\mathbb{N}}$ arise from this correspondence. It is not hard to show that any *ergodic* measure can be obtained in this way; any generic point of the system reflects all the relevant information about the measure, and the desired finite approximations can be read off from such a point. This fact has been noted by a number of authors, including Bergelson, Furstenberg, and Weiss [5, Section 2], Bergelson, Leibman, and Lesigne [6, Section 0.B], Tao [24, Section 2] and Host and Kra [16, Proposition 6.1].

Theorem 3.1 of Section 3 shows that, in fact, *every* shift-invariant measure on $2^{\mathbb{N}}$, ergodic or not, arises in such a way. The proof provides an explicit construction of a sequence of finite combinatorial approximations to any given measure, with, moreover, a uniform upper bound on how large n has to be in order to approximate the measure to a given accuracy.

Section 4 considers some consequences for computable measure theory. There is a precise sense, described in [13, 14, 18, 29], in which a dynamical system can be said to be *computable*; and similarly for a transformation of such a space, and an element of the underlying space. In particular, $2^{\mathbb{N}}$ equipped with the left shift

2010 *Mathematics Subject Classification.* 37A45, 03F60.

Work partially supported by NSF grant DMS-1068829.

is computable, a computable element of $2^{\mathbb{N}}$ is a computable binary sequence, and a computable measure on $2^{\mathbb{N}}$ is an algorithm which computes (arbitrarily good rational approximations to) the measure of each basis set in the usual topology.

It is by now well known that many common ergodic-theoretic constructions are not computable. For example, V'yugin [27, 28] has shown that one cannot generally compute a bound on the rate of convergence of a sequence of ergodic averages $A_n f$, even when f is computable (see also [2, Section 5] and [1]). Similarly, there is a sense in which ergodic decomposition is not computable [17]. The passage from a sequence of sets (A_n) to one of the measures guaranteed to exist by the Furstenberg correspondence principle is certainly not computable, since it is not even continuous in the data.¹ The proof of theorem 3.1 shows, however, that passage in the other direction is fully effective: one can explicitly compute a sequence of combinatorial approximations from the given measure.

This fact has a surprising consequence: every computable shift-invariant measure on $2^{\mathbb{N}}$, ergodic or not, has a computable generic point. A number of recent papers [12, 13, 14] are concerned with identifying conditions under which a computable measure preserving system has such an element. Theorem 4.2 shows that not only is this always the case when the underlying dynamical system is $2^{\mathbb{N}}$ with the shift, but, moreover, one has explicit rates of convergence that are independent of the measure in question.

The existence of generic points for shift-invariant measures on $2^{\mathbb{N}}$ was first established by Colebrook [9]. The construction in Section 4 depends on the fact that one can obtain a point of $2^{\mathbb{N}}$ by piecing together finite specifications, and Sigmund [23] has shown that generic points exist, more generally, for measure-preserving systems satisfying the “specification property.” But it does not seem possible to adapt the computability results here to this more general setting; this is discussed in Section 6.

In recent years, the correspondence principle has been more broadly construed as a way of relating combinatorial configurations in a discrete group with measure-preserving systems on which this group acts. In particular, the principle has been generalized to countable discrete amenable groups in [7], [3, Section 4], and, more recently, [4]. It is noted in [4, Section 1] that one can, conversely, pass in the other direction from “space averages” to combinatorial “group averages” in the case where the action of the group of the space is ergodic. Section 5 below again lifts the restriction to ergodic actions, and provides an effective version of the transformation.

I am grateful to Bryna Kra and Henry Towsner for comments and suggestions on an earlier draft; to Manfred Denker and Matthieu Hoyrup for subsequently pointing me to Sigmund’s results; and, especially, to an anonymous referee for many helpful comments, corrections, suggestions, and references.

2. PRELIMINARIES

For each natural number n , it is convenient to identify n with the set $\{0, \dots, n-1\}$, and to identify each subset A of n with the finite binary string of length n whose i th digit is 1 if and only if i is in A . Note that this representation encodes both the set A and the fact that A is to be viewed as a subset of n . If σ is another binary

¹There is, however, always such a measure that is *low* in the Turing jump of the sequence (A_n) . In this sense the Furstenberg correspondence principle is analogous to the Bolzano-Weierstrass principle; see [20, 21].

sequence, say that σ occurs at position i of A if for every j less than the length of σ , the j th bit of σ agrees with the bit of A at $(i + j) \bmod n$; that is, we let σ wrap around to the beginning of A , if necessary, in doing the comparison. Given A and σ , set

$$\begin{aligned} E_A(\sigma) &= \{i < n \mid \sigma \text{ occurs at position } i \text{ in } A\}, \\ N_A(\sigma) &= |E_A(\sigma)|, \quad \text{and} \\ D_A(\sigma) &= N_A(\sigma)/n. \end{aligned}$$

So $E_A(\sigma)$ is the set of positions at which σ occurs in A , $N_A(\sigma)$ is the number of times it occurs, and $D_A(\sigma)$ is the density of occurrences. Note that $E_A([1]) = A$.

Let $2^{\mathbb{N}}$ denote Cantor space, that is, the space of functions from \mathbb{N} to the discrete space $\{0, 1\}$, under the product topology. If we view elements ω of $2^{\mathbb{N}}$ as infinite binary sequences, it makes sense to write $\sigma \subset \omega$ to denote that the finite sequence σ is an initial segment of ω . The collection of cylinder sets $[\sigma]$ provides a basis for the topology on $2^{\mathbb{N}}$, where $[\sigma] = \{\omega \mid \sigma \subset \omega\}$ is the set of infinite sequences extending σ . I will use \mathcal{B} to denote the Borel sets in this topology. Let T denote the shift-left map on $2^{\mathbb{N}}$, defined by setting $(T\omega)(n) = \omega(n+1)$ for every n . Notice that for every σ , $T^{-1}([\sigma]) = [0\sigma] \cup [1\sigma]$, the set of infinite binary strings in which σ occurs in the second position. Every finite subset A of n gives rise to a measure μ_A on the Borel subsets of Cantor space defined by setting $\mu_A([\sigma]) = D_A(\sigma)$ and applying the Caratheodory extension theorem. By the observations above, we have $\mu_A(T^{-1}[\sigma]) = \mu_A([\sigma])$, which is to say, μ is invariant under the shift.

Now let (A_n) be a sequence of sets of natural numbers, with $A_n \subseteq n$ for each n . By the compactness of the space of measures on Cantor space in the vague topology, there is a subsequence $(\mu_{A_{n_i}})_{i \in \mathbb{N}}$ of (μ_{A_n}) that converges weakly to a measure μ on Cantor space. In particular, for each σ , the sequence $(\mu_{A_{n_i}}([\sigma]))$, which is equal to the sequence $(D_{A_{n_i}}(\sigma))$, converges to $\mu([\sigma])$. Thus we have:

Theorem 2.1. *For every sequence (A_n) of sets with $A_n \subseteq n$, there are a T -invariant measure μ on $2^{\mathbb{N}}$ and a subsequence (A_{n_i}) of (A_n) with the property that for every σ , $\mu([\sigma]) = \lim_{i \rightarrow \infty} D_{A_{n_i}}(\sigma)$.*

This theorem can be proved more directly by iteratively thinning the sequence (A_n) so that the densities converge for each σ , taking a diagonal subsequence, and then defining $\mu([\sigma])$ to be the resulting limit.

We can take Theorem 2.1 to be a precise statement of the Furstenberg correspondence principle, though sometimes the phrase is used to refer to one of its consequences. Note that since the size of σ remains fixed as n grows, the limits in question are not changed if we do not take wraparound into account when counting the number of occurrences of σ in A_n .

3. INVERTING THE CORRESPONDENCE

In this section we complement Theorem 2.1 by showing that, in fact, any shift-invariant measure μ can be obtained as the result of the construction.

Theorem 3.1. *Let μ be any T -invariant measure on $2^{\mathbb{N}}$. Then for each j and $\varepsilon > 0$, there exist $n \leq 2^{O(j/\varepsilon)}$ and $A \subseteq n$ such that for every σ of length at most j , $|\mu([\sigma]) - D_A(\sigma)| < \varepsilon$. Moreover, there is an $m = m(j, \varepsilon)$ such that for any $n \geq m$ there is an $A \subseteq n$ with this property.*

We can abbreviate the conclusion of the theorem by saying that A gives a (j, ε) -good approximation to μ . The first claim provides an explicit bound on how large A needs to be to provide such an approximation. It would be interesting to know whether this bound can be improved. It is the second claim, however, that provides a natural inverse to Theorem 2.1: if for each j we choose m_j large enough to ensure there are $(j, 1/j)$ -good approximations for any $n \geq m_j$, then for any μ we can build a sequence which contains such approximations between m_j to m_{j+1} . Any measure satisfying the conclusion of Theorem 2.1 then has to coincide with μ .

Proof. Fix j and $\varepsilon > 0$, and let k be an integer much larger than j and $1/\varepsilon$, to be specified more precisely later on. Then the set $\{[\tau] \mid \text{length}(\tau) = k\}$ forms a partition of 2^ω , and if $i < k - j$,

$$\mu(T^{-i}[\sigma] \cap [\tau]) = \begin{cases} \mu([\tau]) & \text{if } \sigma \text{ occurs at position } i \text{ of } \tau \\ 0 & \text{otherwise.} \end{cases}$$

Recall that $N_\tau(\sigma)$ denotes the number of times that σ occurs in τ . By the T -invariance of μ , we have

$$\begin{aligned} \mu([\sigma]) &= \frac{1}{k} \sum_{i < k} \mu(T^{-i}[\sigma]) \\ &= \frac{1}{k} \sum_{i < k} \sum_{\tau} \mu(T^{-i}[\sigma] \cap [\tau]) \\ &= \sum_{\tau} \frac{1}{k} \sum_{i < k} \mu(T^{-i}[\sigma] \cap [\tau]) \\ &= \sum_{\tau} \frac{1}{k} (N_\tau(\sigma) + O(j)) \mu([\tau]) \\ &= \sum_{\tau} D_\tau(\sigma) \mu([\tau]) + O(j/k). \end{aligned}$$

In other words, if the τ 's are sufficiently long, the average of the densities of σ in each τ , weighted by $\mu([\tau])$, provide a good approximation to $\mu([\sigma])$. We now obtain the desired set A by concatenating copies of the τ 's in the right proportion; the fact that k is much larger than j will ensure that the occurrences of σ near the border between copies of the τ 's will have a negligible contribution to the overall density.

More precisely, let l be much larger than k and $1/\varepsilon$, and let $\tau_0, \tau_1, \dots, \tau_{2^k-1}$ be an enumeration of the sequences of length k . For each $i \leq 2^k$ let b_i be the closest integer to $(\sum_{m < i} \mu([\tau_m])) \cdot l$, and let $a_{\tau_i} = b_{i+1} - b_i$. Then the values a_τ/l provide a good rational approximation to $\mu([\tau])$, with

$$|\mu([\tau]) - \frac{a_\tau}{l}| < \frac{2}{l} \quad \text{and} \quad \sum_{\tau} a_\tau = b_{2^k} = l.$$

Let A be the set obtained by concatenating a_0 copies of τ_0 , followed by a_1 copies of τ_1 , and so on. Then A has length kl . Accounting for occurrences of σ near the border between such copies, the total number of occurrences of σ in A is given by

$$N_A(\sigma) = \sum_{\tau} a_\tau N_\tau(\sigma) + l \cdot O(j).$$

Dividing by kl , we have

$$\begin{aligned} D_A(\sigma) &= \sum_{\tau} \frac{a_{\tau}}{l} D_{\tau}(\sigma) + O\left(\frac{j}{k}\right) \\ &= \sum_{\tau} \mu([\tau]) D_{\tau}(\sigma) + O\left(\frac{2^k}{l}\right) + O\left(\frac{j}{k}\right) \\ &= \mu([\sigma]) + O\left(\frac{2^k}{l}\right) + O\left(\frac{j}{k}\right), \end{aligned}$$

using the previous expression for $\mu([\sigma])$. Now we only need to choose $k = O(j/\varepsilon)$ large enough to make the second error less than $\varepsilon/2$, and then choose $l = O(2^k/\varepsilon) = 2^{O(j/\varepsilon)}$ to make the first error less than $\varepsilon/2$. The length of A is then $kl = 2^{O(j/\varepsilon)}$, as desired.

The last claim of the proof is easily obtained, for example, by concatenating sufficiently many copies of A and truncating as necessary. \square

As is well known, $2^{\mathbb{N}}$ with the left shift is universal for measure-preserving systems with a distinguished set, in the following sense (see, for example, [25, Example 2.2.6]). Let $\mathcal{X} = (X, \mathcal{C}, \nu, U)$ be a measure-preserving system, and let E be a \mathcal{C} -measurable set. Define a function φ from X to $2^{\mathbb{N}}$ by

$$(\varphi(x))_i = \begin{cases} 1 & \text{if } U^i x \in E \\ 0 & \text{otherwise.} \end{cases}$$

In other words, the 1's in $\varphi(x)$ correspond to places where the orbit of x under U lands in E . Then for any $x \in X$, $T\varphi(x) = \varphi(Ux)$, and $\varphi^{-1}([1]) = E$. Moreover, φ^{-1} , as a function on subsets of $2^{\mathbb{N}}$, is a σ -algebra homomorphism from \mathcal{B} onto the σ -subalgebra of \mathcal{C} generated by E . Define the “push-forward” measure μ on \mathcal{B} by setting $\mu(A) = \nu(\varphi^{-1}(A))$ for every A in \mathcal{B} . Then μ is a T -invariant measure, which encodes information about the measure ν on intersections of finite shifts of E and their complements: if we let $(U^{-i}E)^{\sigma_i}$ denote $U^{-i}E$ if $\sigma_i = 1$ and $X \setminus U^{-i}E$ if $\sigma_i = 0$, we have $\mu([\sigma]) = \nu(\bigcap_{i < \text{length}(\sigma)} (U^{-i}E)^{\sigma_i})$. This allows us to generalize the statement of Theorem 3.1:

Corollary 3.2. *Let $\mathcal{X} = (X, \mathcal{C}, \nu, U)$ be any measure-preserving system, and let E be any \mathcal{C} -measurable set. Then for each j and $\varepsilon > 0$, there exist n and $A \subseteq n$ such that for every σ of length at most j ,*

$$\left| \nu \left(\bigcap_{i < \text{length}(\sigma)} (U^{-i}E)^{\sigma_i} \right) - D_A(\sigma) \right| < \varepsilon.$$

4. CONSEQUENCES FOR COMPUTABLE MEASURE THEORY

We can also consider Theorem 3.1 in computability-theoretic terms. This presupposes some notions from computable analysis and measure theory; I will sketch the necessary background here, and refer the reader to [8, 18, 29] for details.

Computability theory starts with the notion of a *computable* function from the natural numbers to natural numbers, or from finite strings of symbols to finite strings of symbols. One then obtains notions of computability with respect to other finitary objects (integers, pairs of numbers, finite graphs, and so on) by fixing encodings of these objects as numbers or strings. Intuitively, a function from a set of finitary objects to another is said to be computable if there is an algorithm, or

computer program, that computes it. This notion can be made precise using the Turing machine model of computation and fixing the various encodings, but for practical purposes the intuitive description suffices.

Computable analysis has to take into account the representation of infinitary objects, like the real numbers, which cannot be encoded with a finite amount of data. We define a real number r to be computable if there is a computable function f from \mathbb{N} to \mathbb{Q} such that for every i , $|r - f(i)| < 2^{-i}$. In other words, r is computable if one can compute arbitrarily good rational approximations to it. Notice that the choice of $i \mapsto 2^{-i}$ as a rate of convergence is somewhat arbitrary, and the definition is unchanged if one replaces 2^{-i} with any computable sequence of rationals that decreases to 0; from any such representation we can obtain any other. Notice that we have defined a real number to be computable if it has a computable *representation* as a Cauchy sequence of rationals with a fixed rate of convergence; a given computable real will have multiple computable representations.

How shall we define a computable function from \mathbb{R} to \mathbb{R} ? The problem is that the inputs to such a function are no longer finite objects. The standard solution is to say that a function F from \mathbb{R} to \mathbb{R} is computable if there is an algorithm which, on input i , is allowed to ask for rational approximations of the input x to any desired accuracy and, after finitely many such queries, terminates and returns an approximation of $F(x)$ to within 2^{-i} . The notion can be made precise in terms of a Turing machine with access to an oracle tape that contains a representation of the input, but, once again, for practical purposes, the intuitive description suffices. This model of computation on the real numbers was originally proposed by Grzegorzczuk [15], and is an instance of what is generally referred to as “type 2 computability” today [8, 29]. Notice that the algorithm computing F is supposed to act appropriately on any representation of a real number x , whether x is computable or not. One can show that any computable function F from \mathbb{R} to \mathbb{R} is continuous; roughly speaking, this holds because finite approximations to the value of $F(x)$ depend on only a finite amount of the data representing x .

There is nothing special about the real numbers; the method carries over to any system of elements that can be given the structure of a separable metric space. Suppose (X, d) is such a space and A is a countable dense subset of X . Assuming one has finitary representations of the elements of A such that the distances between these elements are computable, then if one replaces \mathbb{R} and \mathbb{Q} in the preceding discussion by X and A , respectively, one obtains notions of a “computable element of X ” and a “computable function on X .” We will not need the full generality of these definitions here. Instead, I will focus on how they play out for $2^{\mathbb{N}}$ and the space of measures on $2^{\mathbb{N}}$.

An element ω of $2^{\mathbb{N}}$ is computable if and only if the function from \mathbb{N} to $\{0, 1\}$ which, on input i , returns the i th digit of ω is computable. Similarly, a computable function T from $2^{\mathbb{N}}$ to $2^{\mathbb{N}}$ is given by an algorithm which, for every i , computes the i th bit of $T\omega$ after querying finitely many bits of ω . For example, if T is the left shift, then T is easily seen to be computable, because in order to output the i th bit of $T\omega$ one need only query the $(i + 1)$ st bit of ω .

A measure μ on $2^{\mathbb{N}}$ is said to be computable if there is an algorithm which, on input σ , computes (arbitrarily good rational approximations to) $\mu([\sigma])$. In other words, μ is computable if there is an algorithm which, on input σ and i , computes a rational approximation of $\mu([\sigma])$ to within 2^{-i} . More generally, one can take

an arbitrary measure μ to be represented by such a function from $S \times \mathbb{N}$ to \mathbb{Q} , where S is the set of finite binary strings. As in the case of the real numbers, it makes sense to talk about algorithms that carry out computations relative to such a representation.

The following theorem provides a sense in which, from a computational point of view, a measure on $2^{\mathbb{N}}$ is “morally equivalent” to a sequence of (j, ε) good approximations.

Theorem 4.1. *There are a computable function $m(j, \varepsilon)$ and an algorithm with the following property: given any representation of a measure μ on $2^{\mathbb{N}}$, the algorithm computes a sequence (A_n) of subsets of n such that for every $n \geq m(j, \varepsilon)$, (A_n) is an (j, ε) -good approximation to μ . Conversely, there is an algorithm which, given a representation of such a function m and sequence (A_n) , computes the measure μ .*

Proof. Let $m(j, \varepsilon)$ be as in the statement of Theorem 3.1. The proof of Theorem 3.1 provides the requisite algorithm, that is, for each n , an explicit description of how to obtain $A \subseteq n$ from finitely many values of μ on cylinder sets. Conversely, given m and (A_n) , to compute $\mu([\sigma])$ to within ε , let $j = \text{length}(\sigma)$, choose $n = m(j, \varepsilon)$, and compute $D_{A_n}(\sigma)$. \square

If μ is a shift-invariant measure on $2^{\mathbb{N}}$, a point ω of $2^{\mathbb{N}}$ is *generic* if for every finite binary sequence σ ,

$$\mu([\sigma]) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i < n} 1_{[\sigma]}(T^i \omega).$$

In other words, for every σ , the limiting frequency of occurrences of σ in ω is $\mu([\sigma])$ (see, for example, [17]). The following theorem shows that given a shift-invariant measure μ on $2^{\mathbb{N}}$, one can compute a generic point, such that the rate of convergence of the limit above is moreover computable (and independent of μ).

Theorem 4.2. *There is a computable function $m(j, \varepsilon)$ with the following property. Given a representation of a shift invariant measure μ on $2^{\mathbb{N}}$, one can compute a point ω that is generic for μ , with the additional property that for every σ of length j , every $\varepsilon > 0$, and every $n \geq m(j, \varepsilon)$, $|\mu([\sigma]) - \frac{1}{n} \sum_{i < n} 1_{[\sigma]}(T^i \omega)| < \varepsilon$.*

Proof. Given μ , for each j let A_j provide a $(j, 2^{-j})$ -good approximation to μ with length bounded as in Theorem 3.1. The idea is to build ω by concatenating copies of A_1 , then copies of A_2 , then copies of A_3 , and so on, choosing enough copies at each stage to ensure that the transitions are smooth. Specifically, construct $\omega = \tau_0 \tau_1 \tau_2 \dots$ in stages, as follows. First, define τ_0 to be the empty sequence. Now, assuming τ_0, \dots, τ_l are defined, set $m_l = \text{length}(\tau_0 \tau_1 \dots \tau_l)$, and let $k = m_l + \text{length}(A_{l+2})$. Let τ_{l+1} be the concatenation of enough copies of A_{l+1} so that $k / \text{length}(\tau_{l+1}) < 2^{-(l+1)}$. Then a routine calculation shows that for every σ of length at most $l+1$ and $n \geq m_{l+1}$, $|\mu([\sigma]) - \frac{1}{n} \sum_{i < n} 1_{[\sigma]}(T^i \omega)| < 2^{-l}$. Clearly ω can be computed from μ , and a bound $m(j, \varepsilon)$ on $m_{\max(j, \lceil \log_2(\varepsilon^{-1}) \rceil + 1)}$ can be computed outright, independent of μ . \square

5. AN EXTENSION TO AMENABLE GROUPS

In recent years the correspondence principle has typically been construed more abstractly as a way of relating combinatorial configurations in a discrete group with measure-preserving systems on which this group acts. The principle has been

generalized to countable discrete amenable groups in [7], [3, Section 4], and even more broadly in [4]. (See also [26].) The conventional way of passing in the other direction, from “space averages” to “group averages,” relies on the pointwise ergodic theorem and works only for ergodic measures. In this section, we provide an effective proof that once again avoids the assumption of ergodicity.

A countable discrete group Γ is said to be *amenable* if for every finite $K \subset \Gamma$ and $\varepsilon > 0$ there is a finite $F \subset \Gamma$ such that $|F \Delta KF| < \varepsilon \cdot |F|$. Given such a Γ , we can fix a sequence $F_0 \subseteq F_1 \subseteq F_2 \subseteq \dots \subseteq \Gamma$ such that $\bigcup_i F_i = \Gamma$ and for every finite set K there is an i such that $|F_j \Delta KF_j| < \varepsilon \cdot |F_j|$ for every $j \geq i$. Such a sequence is called a *Følner sequence*.

Here the natural analogue to $2^{\mathbb{N}}$ is 2^Γ under the product topology. For each $\gamma \in \Gamma$, γ gives rise to the action T_γ on 2^Γ defined by $(T_\gamma \omega)(\alpha) = \omega(\gamma\alpha)$. A measure μ on 2^Γ is said to be Γ -invariant if T_γ preserves μ for each γ . On the natural numbers, $(\{0, \dots, n-1\})_n$ forms a Følner sequence, and it is natural to associate each element of that sequence with the corresponding cyclic subgroup. In general, however, there is no way to associate a group to each element F_n of a Følner sequence, nor a way to paste copies of such groups together. As a result, we need a more general framework.

Fix Γ . Given a finite subset F of Γ and a set X , we define a *partial action of F on X* to consist of a partial function $x \mapsto \gamma x$ for each γ in F , satisfying $1x = x$ for every x , and $\gamma(\gamma'x) = (\gamma\gamma')x$ whenever γ, γ' , and $\gamma\gamma'$ are all in F , and both sides of the equation are defined. Say that the *domain* of F with respect to this partial action is the intersection of the domains of the γ 's, as γ ranges over the elements of F . In other words, an element $i \in X$ is in the domain of the partial action if γi is defined for each i in F .

A pattern, σ , is now a map from some finite subset of Γ to $\{0, 1\}$. As above, the standard topology on 2^Γ is generated by the cylinder sets $[\sigma]$, where $[\sigma] = \{\omega \mid \omega(i) = \sigma(i) \text{ for every } i \in \text{dom}(\sigma)\}$. Fix a finite subset F of Γ and an action of F on some finite set X . If A is a subset of X and i is an element of X , say that σ *occurs at position i in A* if and only if for every α in the domain of σ , $\sigma(\alpha) = 1$ if and only if αi is defined and in A . As in Section 2, define

$$N_A(\sigma) = |\{i \in X \mid \sigma \text{ occurs at position } i \text{ in } A\}|, \quad \text{and} \\ D_A(\sigma) = N_A(\sigma)/n,$$

where the set X and the partial action on X are left implicit in the notation.

The following theorem provides one formulation of the correspondence principle for amenable groups.

Theorem 5.1. *Let Γ be a countable discrete amenable group, with Følner sequence (F_n) . Let (X_n) be a sequence of sets, where each X_n equipped with a partial action of F_n such that $\lim_n |\text{dom}(F_n)|/|X_n| = 1$. Then for any sequence of sets (A_n) , where $A_n \subseteq X_n$ for each n , there are a Γ -invariant measure μ on 2^Γ and a subsequence (A_{n_i}) of (A_n) with the property that for every pattern σ , $\mu([\sigma]) = \lim_{i \rightarrow \infty} D_{A_{n_i}}(\sigma)$.*

Taking $\Gamma = \mathbb{Z}$ and $F_n = \{-(n-1), \dots, n-1\}$ for each n yields a version of Theorem 2.1 with \mathbb{Z} in place of \mathbb{N} . In the formulation in [3, Section 4], for example, the sets X_n are taken to be subsets of Γ itself.

Proof. As in the proof of Theorem 2.1, we can iteratively thin the sequence (A_n) and diagonalize so that the limit in question exists for each σ , and then define

$\mu([\sigma])$ accordingly. We only need to show that μ is additive and Γ invariant, at which point we can apply the Caratheodory extension theorem.

To see that μ is additive, let σ be any pattern, α an element of Γ that is not in the domain of σ , and let σ_0 and σ_1 be the patterns extending σ with value 0 and 1, respectively, at γ . It suffices to show that $\mu([\sigma]) = \mu([\sigma_0]) + \mu([\sigma_1])$. But since (F_n) is a Følner sequence, $\text{dom}(\sigma) \cup \{\gamma\} \subseteq F_n$ for sufficiently large n , and the desired conclusion follows from the fact that $\lim_n |\text{dom}(F_n)|/|X_n| = 1$. Similarly, Γ -invariance also follows from the fact that (F_n) is a Følner sequence with this last property. \square

We have the following inverse:

Theorem 5.2. *Let Γ be a countable discrete amenable group with Følner sequence (F_n) , and let μ be any Γ -invariant measure on 2^Γ . Then for each j and $\varepsilon > 0$, there exist an n , a finite set X , a partial action of F_n on X , and $A \subseteq X$ such that for every σ with domain F_j , $|\mu([\sigma]) - D_A(\sigma)| < \varepsilon$.*

Proof. The proof is similar to that of Theorem 3.1, *mutatis mutandis*. Given j , we can choose k large enough to make $|F_j| \cdot |F_j F_k \Delta F_k|/|F_k|$ arbitrarily small. Then for any σ with domain F_j , if we let τ range over patterns with domain F_k , we have

$$\begin{aligned} \mu([\sigma]) &= \frac{1}{|F_k|} \sum_{i \in F_k} \mu(T_i^{-1}[\sigma]) \\ &= \frac{1}{|F_k|} \sum_{i \in F_k} \sum_{\tau} \mu(T_i^{-1}[\sigma] \cap [\tau]) \\ &= \sum_{\tau} \frac{1}{|F_k|} \sum_{i \in F_k} \mu(T_i^{-1}[\sigma] \cap [\tau]) \end{aligned}$$

Let F_j act partially on F_k by left multiplication, and view τ as representing a subset of F_k . Then as long as i is in the domain of F_j , $T_i^{-1}[\sigma] \cap [\tau]$ is equal either to $[\tau]$ or the empty set, depending on whether σ occurs at position i in τ . But i fails to be in the domain of F_j only when $\gamma i \notin F_k$ for some γ in F_j , and so the set of i that are not in the domain of F_j has cardinality at most $|F_j| \cdot |F_j F_k \Delta F_k|$. Thus we can continue the calculation above,

$$\begin{aligned} \dots &= \sum_{\tau} \frac{1}{|F_k|} (N_{\sigma}(\tau) + O(|F_j| |F_j F_k \Delta F_k|)) \mu([\tau]) \\ &= \sum_{\tau} D_{\sigma}(\tau) \mu([\tau]) + O(|F_j| \cdot |F_j F_k \Delta F_k|/|F_k|). \end{aligned}$$

Now proceed as in the proof of Theorem 3.1. Let X be a disjoint union of copies of F_k , where F_j acts on F_k by left multiplication, insofar as the results of the multiplication land in F_k . Let A be a disjoint union of copies of the various τ 's living on the various F_k 's, where the fraction of occurrences of a given τ approximates $\mu([\tau])$, and $|F_j| \cdot |F_j F_k \Delta F_k|/|F_k|$ is sufficiently small to preserve the quality of the approximation. \square

As in the proof of Corollary 3.2, the universality of 2^Γ means that the result can be pulled back to arbitrary Γ -invariant spaces.

Corollary 5.3. *Let Γ be a countable discrete amenable group with Følner sequence (F_n) . Let $\mathcal{X} = (X, \mathcal{C}, \nu, \Gamma)$ be a measure-preserving system, and let E be any \mathcal{C} -measurable set. Then for each j and $\varepsilon > 0$, there exists a finite set X , a partial action of F_j on X , and $A \subseteq X$ such that for every pattern σ with domain F_j ,*

$$\left| \nu \left(\bigcap_{i \in F_j} (T_i^{-1} E)^{\sigma_i} \right) - D_A(\sigma) \right| < \varepsilon.$$

6. FINAL COMMENTS

The proof of Theorem 3.1 relies on the fact that one can construct a point of $2^{\mathbb{N}}$ by concatenating finite specifications of its orbit behavior. This is an instance of a more general property that some dynamical systems enjoy, known as the *specification property* [23]. (A slightly stronger version is presented in [19].) Let (X, T) be a dynamical system, where X is a compact space and T is a continuous map from X to itself. Let μ be a T -invariant probability measure defined on the Borel subsets of X . A point x in X is said to be generic for μ if for every continuous function f from X to \mathbb{R} , $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i < n} f(T^i x) = \int f d\mu$. The results of Sigmund [23] show that if (X, T) satisfies the specification property, then are generic points for any T -invariant measure, whether it is ergodic or not. The existence of generic points for any shift-invariant measure on $2^{\mathbb{N}}$ is a special case of this result.

The notions of computability discussed in Section 4 can be extended to more general compact metric spaces; see, for example, [13, 14, 18, 29]. In analogy to Theorem 4.2, one might expect that systems satisfying a computable version of the specification property will always have computable generic points. But the methods of Sigmund [23] do not seem to translate to the computable setting: the analogue to Theorem 3.1 above is given by Lemma 1 of Sigmund [22], which relies on the pointwise ergodic theorem in an essential way. In contrast, the proof of Theorem 3.1 relies on particular features of $2^{\mathbb{N}}$. It seems likely that Sigmund's result is noneffective, which is to say, there are computable dynamical systems satisfying a computable version of the specification property, but lacking any computable generic points. It would therefore be interesting to know the extent to which the methods used here can be generalized.

REFERENCES

- [1] Jeremy Avigad. Uncomputably noisy ergodic limits. To appear in the *Notre Dame Journal of Formal Logic*.
- [2] Jeremy Avigad, Philipp Gerhardy, and Henry Towsner. Local stability of ergodic averages. *Trans. Amer. Math. Soc.*, 362(1):261–288, 2010.
- [3] Vitaly Bergelson. Ergodic theory and Diophantine problems. In F. Blanchard, A. Maass, and A. Nogueira, editors, *Topics in symbolic dynamics and applications*, pages 167–205. Cambridge Univ. Press, Cambridge, 2000.
- [4] Vitaly Bergelson and Hillel Furstenberg. WM groups and Ramsey theory. *Topology Appl.*, 156(16):2572–2580, 2009.
- [5] Vitaly Bergelson, Hillel Furstenberg, and Benjamin Weiss. Piecewise-Bohr sets of integers and combinatorial number theory. In M. Klazar et al., editors, *Topics in discrete mathematics*, pages 13–37. Springer, Berlin, 2006.
- [6] Vitaly Bergelson, Alexander Leibman, and Emmanuel Lesigne. Complexities of finite families of polynomials, Weyl systems, and constructions in combinatorial number theory. *J. Anal. Math.*, 103:47–92, 2007.

- [7] Vitaly Bergelson and Randall McCutcheon. Recurrence for semigroup actions and a non-commutative Schur theorem. In M. G. Nerurkar, D. P. Dokken and D. B. Ellis., *Topological dynamics and applications*, pages 205–222. Amer. Math. Soc., Providence, RI, 1998.
- [8] Vasco Brattka, Peter Hertling, and Klaus Weihrauch. A tutorial on computable analysis. In S. Barry Cooper, Benedikt Löwe, and Andrea Sorbi, editors, *New computational paradigms: changing conceptions of what is computable*, pages 425–491. Springer, New York, 2008.
- [9] C. M. Colebrook. The Hausdorff dimension of certain sets of nonnormal numbers. *Michigan Math. J.*, 17:103–116, 1970.
- [10] H. Furstenberg. Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Analyse Math.*, 31:204–256, 1977.
- [11] H. Furstenberg. *Recurrence in ergodic theory and combinatorial number theory*. Princeton University Press, Princeton, N.J., 1981.
- [12] Stefano Galatolo, Mathieu Hoyrup, and Cristóbal Rojas. A constructive Borel-Cantelli lemma. Constructing orbits with required statistical properties. *Theor. Comput. Sci.*, 410(21-23):2207–2222, 2009.
- [13] Stefano Galatolo, Mathieu Hoyrup, and Cristóbal Rojas. Computing the speed of convergence of ergodic averages and pseudorandom points in computable dynamical systems. To appear in the *Proceedings of Computability and Complexity in Analysis (CCA) 2010*, arXiv:1006.0392.
- [14] Stefano Galatolo, Mathieu Hoyrup, and Cristóbal Rojas. Dynamical systems, simulation, abstract computation. arXiv:1101.0833.
- [15] Andrzej Grzegorzczak. On the definitions of computable real continuous functions. *Fundamenta Mathematicae*, 44:61–71, 1957.
- [16] Bernard Host and Bryna Kra. Uniformity seminorms on ℓ^∞ and applications. *J. Anal. Math.*, 108:219–276, 2009.
- [17] Mathieu Hoyrup. Randomness and the ergodic decomposition. In Benedikt Löwe et al., editors, *Computability in Europe (CiE) 2011*, pages 122–131, Springer, Berlin, 2011.
- [18] Mathieu Hoyrup and Cristóbal Rojas. Computability of probability measures and Martin-Löf randomness over metric spaces. *Inform. and Comput.*, 207(7):830–847, 2009.
- [19] Anatole Katok and Boris Hasselblatt. *Introduction to the modern theory of dynamical systems*. Cambridge Univ. Press, Cambridge, 1995.
- [20] Alexander P. Kreuzer. The cohesive principle and the Bolzano-Weierstrass principle. *Math. Log. Q.*, 57(3):292–298, 2011.
- [21] Pavol Safarik and Ulrich Kohlenbach. On the computational content of the Bolzano-Weierstraß principle. *Math. Log. Q.*, 56(5):508–532, 2010.
- [22] Karl Sigmund. Generic properties of invariant measures for Axiom A diffeomorphisms. *Invent. Math.*, 11:99–109, 1970.
- [23] Karl Sigmund. On dynamical systems with the specification property. *Trans. Amer. Math. Soc.*, 190:285–299, 1974.
- [24] Terence Tao. Norm convergence of multiple ergodic averages for commuting transformations. *Ergodic Theory Dynam. Systems*, 28(2):657–688, 2008.
- [25] Terence Tao. *Poincaré’s legacies, part I: pages from year two of a mathematical blog*. Amer. Math. Soc., Providence, RI, 2009.
- [26] Henry Townser. A general correspondence between averages and integrals. arXiv:0804.2773, 2008.
- [27] V. V. V’yugin. Ergodic convergence in probability, and an ergodic theorem for individual random sequences. *Teor. Veroyatnost. i Primenen.*, 42(1):35–50, 1997.
- [28] V. V. V’yugin. Ergodic theorems for individual random sequences. *Theoret. Comput. Sci.*, 207(2):343–361, 1998.
- [29] Klaus Weihrauch. Computability on the probability measures on the Borel sets of the unit interval. *Theoret. Comput. Sci.*, 219(1-2):421–437, 1999.

DEPARTMENTS OF PHILOSOPHY AND MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA 15213